Linux Database HOWTO (PostgreSQL Relational Database System)

Al Dev (Alavoor Vasudevan) deepav@pol.net v1.0, 19 July 1997

This document is about HOW-TO setup a International Standard/ANSI Relational SQL Database "PostgreSQL" on your unix system which can be used as Application Database Server or a Web Database Server. This document also gives information on the interface programs for the database like Front End GUIs, RAD tools (Rapid Application Development), programming languages interfaces ("C", "C++", Java, Perl), ODBC, JDBC drivers and Web Database Tools and Interface programs. Information given here applies to all other unix platforms. The information given hereby will be very useful for persons who are new to PostgreSQL, Databases and SQL language.

Contents

1 Introduction **7**

2 What is PostgreSQL ? **7**

3 Where to get it? **8**

4 Regression Test Package to validate PostgreSQL against ISO/ANSI SQL Standards **8**

5 GUI FrontEnd Tool for PostgreSQL (Graphical User Interface) **9**

6 ODBC Drivers for PostgreSQL **9**

7 UDBC Drivers for PostgreSQL **9**

8 JDBC Drivers for PostgreSQL **9**

9 Kanchenjunga - Java RAD Tool for PostgreSQL **10**

10 Java Classes for PostgreSQL **10**

11 Perl 5 interface for PostgreSQL **10**

12 Windows Interactive Query Tool for PostgreSQL (WISQL) **10**

13 Interactive Query Tool - ISQL for PostgreSQL **11**

14 Problem/Project Tracking System Application Tool for PostgreSQL **11**

15 PostgreSQL 4GL for web database applications - AppGEN Development System **11**

16 Web Database Design/Implementation tool for PostgreSQL - EARP **12**

16.1 What is EARP? 12

16.2 Implementation 12

16.3 What you need to run EARP 13

16.4 News Flash 13

16.5 How does it work? 13

16.6 Some online examples 14

16.7 Where do I get it? 14

16.8 Available Documentation 14

16.9 A History of EARP 14

17 WWW Web interface for PostgresSQL - dbengine **15**

18 Apache Webserver Module for PostgreSQL - NeoSoft NeoWebScript **16**

19 PHP/FI Server-side html-embedded scripting language for PostgreSQL **17**

19.1 Major Features 17

19.2 Credits 19

19.3 Brief History 19

19.4 So, what can I do with PHP/FI? 20

19.5 A simple example 20

19.6 CGI Redirection 21

19.7 Running PHP/FI from the command line 22

20 Python Interface for PostgreSQL **22**

20.1 Where to get ... ? 23

20.2 Information and support 23

21 Gateway between PostgreSQL and the WWW - WDB-P95 **24**

21.1 Does the Postgres95 server,pgperl, and httpd have to be on the same host? 24

21.2 New Version 25

22 "C" language Interface for PostgreSQL **25**

23 "C++" language Interface for PostgreSQL **25**

24 ESQL/C for PostgreSQL **25**

25 Japanese Kanji Code for PostgreSQL **25**

26 PostgreSQL Port to Windows 95/Windows NT **26**

27 Mailing Lists **26**

28 Documentations and Books **27**

29 Technical support for PostgreSQL **27**

30 Economic and Business Aspects **27**

31 Conclusion **28**

32 FAQ - Questions on PostgreSQL **28**

32.1 Frequently Asked Questions (FAQ) for PostgreSQL 28
32.2 What is PostgreSQL? 29
32.3 What does PostgreSQL run on? 29
32.4 Where can I get PostgreSQL? 30
32.5 What's the copyright on PostgreSQL? 30
32.6 Support for PostgreSQL 30
32.7 Latest release of PostgreSQL 31
32.8 Is there a commercial version of PostgreSQL? 31
32.9 What documentation is available for PostgreSQL? 31
32.10What version of SQL does PostgreSQL use? 31
32.11Does PostgreSQL work with databases from earlier versions of postgres? 32
32.12How many people use PostgreSQL? 32

33 FAQ - Installation Questions **32**

33.1 Initdb doesn't run 32
33.2 When I start up the postmaster, I get "FindBackend- could not find a backend to execute..." "postmaster- could not find backend to execute..." 32

33.3 The system seems to be confused about commas, decimal points, and date formats. 32
33.4 How do I install PostgreSQL somewhere other than /usr/local/pgsql? 33
33.5 When I run postmaster, I get a Bad System Call core dumped message. 33
33.6 When I try to start the postmaster, I get IpcMemoryCreate errors. 33
33.7 I have changed a source file, but a recompile does not see the change? 33

34 FAQ - PostgreSQL Features **33**
34.1 How do I specify a KEY or other constraints on a column? 33
34.2 Does PostgreSQL support nested subqueries? 33
34.3 How do I define a unique indices? 33
34.4 I've having a lot of problems using rules. 34
34.5 I can't seem to write into the middle of large objects reliably. 34
34.6 Does PostgreSQL have a graphical user interface? A report generator? A embedded query language interface? 34
34.7 How can I write client applications to PostgreSQL? 34
34.8 How do I prevent other hosts from accessing my PostgreSQL backend? 34
34.9 How do I set up a pg_group? 34
34.10What is the exact difference between binary cursors and normal cursors? 35
34.11Why doesn't the != operator work? 35
34.12What is a R-tree index and what is it used for? 35
34.13What is the maximum size for a tuple? 36
34.14I defined indices but my queries don't seem to make use of them. Why? 36
34.15Are there ODBC drivers for PostgreSQL? 36
34.16How do I use postgres for multi-dimensional indexing (¿ 2 dimensions)? 36
34.17How do I do regular expression searches? case-insensitive regexp searching? 37
34.18I can't access the database as the 'root' user. 37
34.19I experienced a server crash during a vacuum. How do I remove the lock file? 37
34.20What is the difference between the various character types? 37
34.21In a query, how do I detect if a field is NULL? 38
34.22How do I see how the query optimizer is evaluating my query? 38
34.23How do I create a serial field? 38
34.24How do I create a multi-column index? 39
34.25What are the temp_XXX files in my database directory? 39

34.26Why are my table files not getting any smaller after a delete? 39

34.27Why can't I connect to my database from another machine? 39

34.28I get the error 'default index class unsupported' when creating an index. How do I do it? . . 39

34.29Why does creating an index crash the backend server? 39

34.30How do I find out what indexes or operations are defined in the database? 40

34.31Why do statements require an extra character at the end? Why does 'createuser' return 'unexpected last match in input()'? Why does pg_dump fail? 40

34.32All my servers crash under concurrent table access. Why? 40

34.33What tools are available for hooking postgres to Web pages? 40

34.34What is the time-warp feature and how does it relate to vacuum? 40

34.35How do I tune the database engine for better performance? 41

34.36What debugging features are available in PostgreSQL? 41

34.37What is an oid? What is a tid? 42

34.38What is the meaning of some of the terms used in Postgres? 42

34.39What is Genetic Query Optimization? 42

34.40I am running Solaris and my dates display wrong. Why? 42

34.41How do I enable more than 32 concurrent backends? 42

35 FAQ - Extending PostgreSQL 43

35.1 I wrote a user-defined function and when I run it in psql, it dumps core. 43

35.2 I get messages of the type NOTICE-PortalHeapMemoryFree- 0x402251d0 not in alloc set! . . 43

35.3 I've written some nifty new types and functions for PostgreSQL. 43

35.4 How do I write a C function to return a tuple? 43

36 FAQ - Bugs 43

36.1 How do I make a bug report? 43

37 FAQ Linux-PostgreSQL - Compiling PostgreSQL 44

37.1 What changes do I need to make to src/Makefile.global or src/Makefile.custom? 44

37.2 Why do I get problems with missing libreadline? 44

37.3 *REDHAT* Why do I get problems with missing libdl and dlfcn.h? 44

37.4 *SLACKWARE*3.1 Why do I get problems with missing libdl and dlfcn.h? 45

37.5 My compile of the backend dies complaining about the include file dlfcn.h missing 46

37.6 GCC complains about an ignored option -fpic 46

37.7 I get warnings of the form warning- cast from pointer to integer of different size 46

37.8 $SuSE - Linux$4.2 Where is curses and termcap? 46

37.9 Why do I get problems with ld.so? 46

37.10Why do I get 'yy_flush_buffer undefined' errors? 47

37.11How do I compile PostgreSQL on an a.out system? 47

37.12Why does make fail with - yacc -d /disk2/PostgreSQL/src/backend/parser/gram.y make- /usr/bin/make- cannot execute binary file 47

37.13What are the references in X11_LIB to libsocket and libnsl in src/Makefile.global? 47

37.14$DEBIAN$ Where is libtermcap? 47

38 FAQ Linux-PostgreSQL - Compiling accessory programs 47

38.1 The linker fails to find libX11 when compiling pgtclsh 47

39 FAQ Linux-PostgreSQL - Runtime Problems 48

39.1 I get an error reporting _fUnKy_POSTPORT_sTuFf_ undefined when running scripts like createuser 48

39.2 I run postmaster and after that system says 'Bad system call(Core dumped)' 48

39.3 When I try to start the Postmaster, why do I get an error of the form Failed Assertion("!(file != 0) - (null)", File 48

39.4 Why doesn't createuser work? 48

39.5 Why do I get an error like- IpcMemoryCreate- memKey=155356396 , size=760632 , ... 48

39.6 Why does psql fail with- psql- can't load library 'libpq.so.1' 48

40 FAQ IRIX-PostgreSQL - Compiling PostgreSQL 49

41 FAQ-IRIX Installing Postgres95 49

41.1 What extra items do I need to install Postgres95 under Irix? 49

41.2 What changes do I need to make to src/Makefile.global or src/Makefile.custom? 50

41.3 What are the references in X11_LIB to libsocket and libnsl in src/Makefile.global? 50

41.4 Are there any other changes I should make? 50

41.5 Can I install PostgreSQL under Irix 6.4? 50

42 FAQ-IRIX Deinstalling Postgres95 51

42.1 Why can't I move the executable files? 51

43 FAQ-IRIX Extending Postgres95 **51**

43.1 How do I compile a C program to create a function for extending Postgres95 51

44 FAQ-IRIX PostgreSQL AUTHOR contacts **51**

45 Copyright Notice **51**

1 Introduction

Every computer system needs a database to store/retrieve the information. The primary reason for the invention of computer is to store, retrieve and process information and do all these very quickly. The most popular database systems are based on the International Standard Organisation (ISO) SQL specifications which are also based on ANSI SQL (american) standards. Current specifications generally used are ANSI SQL 92 and ANSI SQL 89. Popular database like Oracle, Sybase and Informix systems are based on these standards.

The world's most popular FREE Database which is ISO SQL, ANSI SQL/92 and ANSI SQL/89 compliant RDBMS is PostgreSQL, also known as pgsql (pronounced as "pee-gee-sequel"). PostgreSQL is the only free RDBMS in the world which supports ANSI SQL standards. This document will tell you how-to install the database and also all the packages related to databases, how to set up the Web database, application database, front end GUIs and interface programs. It is strongly advised that you MUST write your database applications 100 % compliant to standards of ISO/ANSI SQL, ODBC, JDBC so that your application is portable across multiple databases like PostgreSQL, Oracle, Sybase, Informix etc.

2 What is PostgreSQL ?

PostgreSQL Version 6.1.1 is free database, complete source code is given to you and is a Relational Database System compliant with ANSI SQL92 and SQL89 and runs on diverse hardware platforms and Operating systems. Millions of PostgreSQL is installed as Data servers, Datacenter servers, Web databases, WWW http database servers and Application servers. It is much more advanced and is a object oriented relational database (ORDBMS). PostgreSQL can store more data types than traditional datatypes like integer, characters, images, etc. - it can store complex object datatypes like objects and classes (Version 7.0 will make this capabilities more advanced). PostgreSQL runs on Solaris, SunOS, HPUX, AIX, Linux, Irix, Digital Unix, BSDi,NetBSD, FreeBSD, SCO unix, NEXTSTEP, Unixware and all and every flavor of Unix. Port to Windows 95/NT is underway.

- Title: PostgreSQL ANSI SQL92/ANSI SQL89 RDBMS Database (Relational Database Management System)
- Current Version: 6.1.1
- Age: PostgreSQL is 12 years old. Developed since 1985
- Authors: Developed by millions/universities/companies on internet for the past 12 YEARS

3 Where to get it?

The primary Web site is given as below :

- "Binaries only" distribution of PostgreSQL : You can run PostgreSQL without compiling the source. Get binaries for Intel-Linux from `<http://www.redhat.com/pub/contrib/i386/>` file is postgresql-6.1-6.i386.rpm. This is in the redhat package 'rpm' format and it contains both source and binaries for PostgreSQL.
- Primary Web site: `<http://www.postgresql.org/>`
- Secondary Web site: `<http://logical.thought.net/postgres95/>`
- `<http://www.itm.tu-clausthal.de/mirrors/postgres95/>`
- `<http://s2k-ftp.cs.berkeley.edu:8000/postgres95/>`

The ftp sites are listed below :-

- "Binaries only" distribution of PostgreSQL : You can run PostgreSQL without compiling the source. Get binaries for Intel-Linux from `<ftp://ftp.redhat.com/pub/contrib/i386/>` file is postgresql-6.1-6.i386.rpm. This is in the redhat package 'rpm' format and it contains both source and binaries for PostgreSQL.
- Primary FTP: `<ftp://ftp.postgresql.org/pub>`
- Secondary FTP: `<ftp://ftp.chicks.net/pub/postgresql>`
- `<ftp://ftp.emsi.priv.at/pub/postgres/>`
- `<ftp://ftp.itm.tu-clausthal.de/pub/mirrors/postgres95>`
- `<ftp://rocker.sch.bme.hu/pub/mirrors/postgreSQL>`
- `<ftp://ftp.jaist.ac.jp/pub/dbms/postgres95>`
- `<ftp://ftp.luga.or.at/pub/postgres95>`
- `<ftp://postgres95.vnet.net:/pub/postgres95>`

4 Regression Test Package to validate PostgreSQL against ISO/ANSI SQL Standards

Regression test package is included in the distribution and let your computer do the verfication for you to check if all the ANSI SQL 92 /ISO SQL commands run as per technical specifications. It is better to check always, but never guess! Advantage is that computer can rapidly test thousands of SQL tests in a very short time. The test package already contains hundreds of SQL test programs. You can add more tests just in case you need to, and can upload to the primary PostgreSQL web site if you feel that will be useful for others on internet.

5 GUI FrontEnd Tool for PostgreSQL (Graphical User Interface)

PostgreSQL has TCL/TK interface library in the distribution called 'pgtcl'. TCL/TK is a Rapid Application Development tool and is a very powerful scripting language. Develop once and run it everywhere on NT, Win 95, Linux and all unixes! TCL/TK is also widely used as a Internet script language. So you will have only one language for all your needs - applications and internet. TCL stands for 'Tool Command Language' and TK is 'Tool Kit'. TCL/TK is usually shipped with every linux cdrom. Also you can get it from these sites -

- `<http://sunscript.sun.com/>`
- `<http://sunscript.sun.com/TclTkCore/>`
- `<ftp://ftp.sunlabs.com/pub/tcl/tcl8.0a2.tar.Z>`
- Reference text book: Many textbooks on TCL/TK are available in the market.

6 ODBC Drivers for PostgreSQL

ODBC stands for 'Open DataBase Connectivity' is a popular standard for accessing information from various databases from different vendors. Applications written using the ODBC drivers are guaranteed to work with various databases like PostgreSQL, Oracle, Sybase, Informix etc..

- `<http://www.ids.net/ bjepson/freeODBC/>` This is a cost free version of ODBC.
- `<http://www.openlinksw.com>` Open Link Software Corporation is selling ODBC for PostgreSQL and other databases.

7 UDBC Drivers for PostgreSQL

UDBC is a static version of ODBC independent of driver managers and DLL support, used to embed database connectivity support directly into applications.

- `<http://www.openlinksw.com>` Open Link Software Corporation is selling UDBC for PostgreSQL and other databases.

8 JDBC Drivers for PostgreSQL

JDBC stands for 'Java DataBase Connectivity'. Java is a platform independent programming language developed by Sun Microsystems. Java programmers are encouraged to write database applications using the JDBC to facilitate portability across databases like PostgreSQL, Oracle, informix, etc. If you write Java applications you can get JDBC drivers for PostgreSQL from the following sites:

- `<http://www.demon.co.uk/finder/postgres/index.html>` Sun's Java connectivity to PostgreSQL
- `<ftp://ftp.ai.mit.edu/people/rst/rst-jdbc.tar.gz>`
- `<http://www.openlinksw.com>` Open Link Software Corporation is selling JDBC for PostgreSQL and other databases.

9 Kanchenjunga - Java RAD Tool for PostgreSQL

Kanchenjunga is a Java Rapid Application Tool for PostgreSQL. You can use this tool to develop rapidly the java application interfacing PostgreSQL.

- `<http://www.man.ac.uk/ whaley/kj/kanch.html>`

10 Java Classes for PostgreSQL

Java programmers can find these classes for PostgreSQL very useful.

- `<ftp://www.blackdown.org/pub/Java/Java-Postgres95>`
- `<http://www.blackdown.org>`

11 Perl 5 interface for PostgreSQL

It included in the distribution of PostgreSQL. Check in src/pgsql_perl5 directory.

- Email: E.Mergl@bawue.de
- Another source from -
- Perl Home page :

12 Windows Interactive Query Tool for PostgreSQL (WISQL)

Identical to Microsoft SQL server WISQL! It has nice GUI and has history of commands. Also you can cut and paste and it has other nice features to improve your productivity.

- `<http://www.troubador.com/ keidav/index.html>`
- Email: keidav@whidbey.com
- `<http://www.ucolick.org/ de/>` in file tcl_syb/wisql.html
- `<http://www.troubador.com/ keidav/index.html>`
- Email: de@ucolick.org

13 Interactive Query Tool - ISQL for PostgreSQL

ISQL is For Character command line terminals. This is included in the distribution, and is called PSQL. Very similar to Sybase ISQL, Oracle SQLplus. At unix prompt give command 'psql' which will put you in psql¿ prompt. Type /h to see help of commands. Very user friendly and easy to use. Also very useful for shell scripting in Bourne, Korn and C-shells.

14 Problem/Project Tracking System Application Tool for PostgreSQL

This is at

- `<http://www.homeport.org/ shevett/pts/>`

15 PostgreSQL 4GL for web database applications - AppGEN Development System

AppGEN can be downloaded from

`<http://www.man.ac.uk/ whaley/ag/appgen.html>`

`<ftp://ftp.mcc.ac.uk/pub/linux/ALPHA/AppGEN>`.

Extract from the home page of AppGEN is given below:-

AppGEN is a high level fourth generation language and application generator for producing World Wide Web (WWW) based applications. These applications are typically used over the internet or within a corporate intranet. AppGEN applications are implemented as C scripts conforming to the Common Gateway Interface (CGI) standard supported by most Web Servers.

To use AppGEN you will need the following :-

PostgresSQL, relational database management system

A CGI compatible web server such as NCSA's HTTPD

An ansi C compiler such as GCC

AppGEN consists of the following Unix (Linux) executables :-

- defgen, which produces a basic template application from a logical data structure. The applications are capable of adding, updating, deleting and searching for records within the database whilst automatically maintaining referential integrity.

- appgen, the AppGEN compiler which compiles the appgen source code into CGI executable C source and HTML formatted documents ready for deployment on a web server.

- dbf2sql, a utility fo converting dBase III compatible .dbf files into executable SQL scripts. This enables data stored in most DOS/Windows based database packages to be migrated to a SQL server such as PostgresSQL.
- In addition, AppGEN comprises of a collection of HTML documents, GIF files and Java applets which are used at runtime by the system. And of course, like all good software, the full source code is included.

The author, Andrew Whaley, can be contacted on andrew@arthur.smuht.nwest.nhs.uk and would appreciate any comments or suggestions about the software.

16 Web Database Design/Implementation tool for PostgreSQL - EARP

`<http://www.oswego.edu/Earp>`

`<ftp://ftp.oswego.edu>` in the directory 'pub/unix/earp'.

The extract from the home page of EARP is given below:-

The "Easily Adjustable Response Program" (EARP) created by David Dougherty

16.1 What is EARP?

EARP is a Web Database Design/Implementation tool, built on top of the Postgres95 database system. Its functionality includes:

- A Visual Design System.
- A sendmail interface. (can handle incoming and outgoing mail)
- An Enhanced Security Mechanism.
- A cgi driver.

16.2 Implementation

The main implementation of EARP is a CGI binary which runs under the http daemon to provide access to the database server. All of the design tools are built into the driver, no design takes place over anything but the web. The tools themselves require a graphical browser, the compatability of objects designed with the tools is implementation independent, based on designing individuals preferences.

16.3 What you need to run EARP

EARP will likely run on a variety of platforms with little or no porting. The known working platforms consist of the following:

- Solaris 2.5
- Linux 1.2.13+
- GNU C++
- Postgres95 (Version 1.01 / 1.02)
- netsite server
- NCSA httpd

- GNU C++
- Postgres95 (Version 1.01 / 1.02)
- NCSA httpd
- Apache httpd

16.4 News Flash

The current (1.3) release of Earp was designed on top of the libpq release that came with Postgres95 v1.01/1.02. If you are using a more recent version of Postgres, expect that the program will require some porting to work correctly. In the development version (Earp 2.0), libpq support is being incorperated as a module, and thus will support as many versions of postgres as we have time to write the modules. The development release is expected to become public near mid-spring(97).

16.5 How does it work?

One of the main features of EARP is that it uses an Object Oriented approach to producing html pages which interface to the database. Most pages will consist of several objects. Each object is produced by some sort of tool and given a name, objects are then linked together in a callable sequence by the page tool. Objects are also reusable across multiple pages. Basic tools exist for HTML, Querys, Grabbing input from forms, Extendable Formatting of Query and Input objects, and Linking together of objects into other objects. More advanced tools include the mail tool and the multithreaded query tool.

Another feature of EARP is advanced security. Access to various areas of the EARP system can be limited in a variety of ways. To facilitate its advanced security, EARP performs checks for each connection to the system, determining what ids and groups the connecting agent belongs to. Access to areas is defined seperately, and the combination decides if access to a specific area of Earp is allowed. Moreover, all that is required to implement the security features is an http server that supports basic (or better) user authentication.

16.6 Some online examples

As part of the ICC Help Database, the Catalog Search Page is an EARP document which runs several queries. The selection boxes are generated by the EARP program from listings in the database.

As another example of what can be done using EARP... now you can look at the List of Objects in the Help Database.

Creating the three interfaces for the link took me less than 15 minutes.

16.7 Where do I get it?

EARP is available via anonymous ftp from `<ftp://ftp.oswego.edu>` in the directory 'pub/unix/earp'. The version as of this writing is 1.3.1

Please, once you've retrieved EARP and gotten it to work, drop me a line and tell me your success or failure story.

16.8 Available Documentation

All documentation has been moved to the User Docs and Tutorials index page.

16.9 A History of EARP

Earp 0.1 began in Fall of 1995 as program I was working on to build a dynamically configurable web accessable guest book. At that point it was a bunch of cgi programs that all did different but usefull things and were held together with SSI glue, and a little sneaky c programming. What I soon realized though is that I was doing a lot of repetitive work, and that most of what I was doing had to be run in many windows at once (netscape, emacs, shell, mail) for it to make any sense, and that debugging was quickly becomming a nightmare. At that time I was also being approached by my friend and boss Don Michaels, who was interested in automating a large hunk of our user support, and keeping a historical database of requests and responses.

Soon, I had worked out the initial scheme for what is now quickly becomming our help database, only I balked at the idea of building a help database with what was at that point a very primitive set of utilities. When spring classes were occuring(96) I started it anyway, mainly out of boredom, but also because I was in a database design class and a wanted to flex a few brain muscles. After a while I had a reasonable prototype up and running, which made Don very happy as he had basically given up on the idea of anyone every really creating a help database for him. (The protoype is still running on one of my servers...(june96)) The prototype did some very interesting things, but by april I was again getting discouraged... Everytime I wanted to change something, I had to go through a lengthy process of recompilation, or find an entry in a text file full of distractions. Also, there was no way for me to use the building block idea which is so usefull in EARP... I did a lot of huge cutting and pasting. About the time that classes where ending I had again given up on the current scheme of things, and decided that what I needed was a better set of tools for what I was doing. Additionally, I also wanted to make my prototype work on top of a REAL realtional database, and I wasn't cherishing the idea of reworking all those hard coded accesses, links, and output methods.

I had a break for a short while, if you want to call it that. We sponsored the SUNY CIT Conference and I was so busy for about a week and a half that I got distnaced from most of what I had written for the prototype help database, except the ideas I had had when I wrote the initial series of utilities, and what my biggest peeves were about the current state of things.

Shortly after the conference, I began the prototype for the current version of EARP(may96), using the postgres95 database as my relational backend. By the middle of June, the prototype had evolved into a fairly nice integrated suite of tool prototypes, with the primary advantage that they all ran over html, and stored their initialization information in the database. Most of the second half of june was spent debugging and working out the kinks in the code, and playing with the interface.(within those two weeks I accessed our web server over 5000 times.) By the end of June I had most of the major bugs blasted out of EARP, and a large enough number of objects in the "new" help database to officially announce it to our help support staff.

Incidentally, I also kept a journal during the development of EARP, and myself and Don Michaels are presenting a paper describing the Help Database at the Chicago ACM/SIGUCCS conference in September of this year.

17 WWW Web interface for PostgresSQL - dbengine

`<http://www.cis-computer.com/dbengine/ >`

The extract from the home page of dbengine is given below:-

dbengine a plug 'n play Web interface for Postgres95 created by Ingo Ciechowski

Version 0.82 alpha Documentation as of 11/23/96

About dbengine

dbengine is an interface between the WWW and Postgres95 which provides simple access to any existing database within just a few minues.

This little Perl program was born after I've tried quite a lot of already available packages like AppGen, , PHP-FI and more. So why did I kind of re-invent the wheel ?

Well, PHP-FI gives you a Perl like language in your documents, but no real Perl while AppGen and wdb-p95 require that you create some configuration file for each of your databases – sound's like you'll first of all have to learn some sort of new mata language before you can get started.

That was the point when I started to feel more familiar with a small Perl applet of my own... and now my dbengine became a tool that I think at is ready to be shared with others.

Unlike other tools you don't have to learn any special programming or scripting language to get started with dbengine. Also there's no configuration file for each database, so you don't have to get familiar with such a new structure. However - in case you want to gain access to the full features of dbengine it'd be a good idea to know the Perl language.

The whole system can be configured by simple manipulations of an additional database that contains closer information about how to visualize your database access. You can even specify virtual Fields which are calculated on the fly right before they're displayed on the screen.

License

dbengine is free software under the same terms as Perl. Read its licence if you aren't sure what you can or can't do. The bottom line is that this is a kinder and gentler version of the GNU licence – one that doesn't infect your work if you care to borrow from dbengine or package up pieces of it as part of a commercial product!

18 Apache Webserver Module for PostgreSQL - NeoSoft NeoWebScript

Apache is a well-known Web Server. And a module to interface PostgreSQL to Apache Webserver is at - `<http://www.neosoft.com/neowebscript/>`

The extract from the home page of NeoWebScript is given below:-

NeoWebScript is a programming language that allows both simple and complex programs to be embedded into HTML files.

When an HTML page containing embedded NeoWebScript is requested, the NeoWebScript-enabled webserver executes the embedded script(s), producing a webpage containing customized content created by the program.

NeoWebScript is a fast, secure, easy to learn way to do powerful, server-based interactive programming directly in the HTML code in web pages. With NeoWebScript, counters, email forms, graffiti walls, guest books and visitor tracking are all easy, even for a beginning programmer. See how well NeoWebScript holds its' own vs. PERL and JavaScript.

NeoWebScript 2.2 just released! On June 24, 1997, NeoSoft released NeoWebScript 2.2, integrating it with the new Apache 1.2.0 server.

If you're on a webserver that has NeoWebScript installed and you would like to started with it, we have a lot of User Info available. Our New User FAQ has the basic answers to get you started. The Tutorials guide you through learning the language, while the Demos give you prebuilt applications you can download and modify. Commands and Variables are the complete language references, and Troubleshooting contains hints and tips to help you get past any problems.

If you'd like to install NeoWebScript on your webserver, your Webmaster needs to read our Sysop FAQ to get started. Theory of Operations will explain how NeoWebScript works, while installation will take them through the steps. Management deals with configuration issues and running the server, tests let you verify correct NeoWebScript operation, and troubleshooting deals with server problems.

Hey, wait a minute you ask, how much do you want for all this great software, huh? There is no cost to you to use NeoWebScript-2.2 for your ISP, your intranet, or your extranet. You'll see a full license when you register to download, but the gist is we'd like a whopping $ 99 if you want to embed it in your own product or use it in a commerce (eg. SSL) server.

NeoWebScript is a module for the Apache webserver that allows you to embed the Tcl/Tk programming language in your webpages as a scripting tool. It was invented by Karl Lehenbauer, NeoSoft's Chief Technical Officer, and documented, enhanced and extended by NeoSoft's programmers and technical writers.

The Apache webserver is the world's most popular webserver, accounting for 42 % of the 1,044,163 sites polled by the May 1997 Netcraft Web Server survey. The next largest entry were the various Microsoft servers, reporting in with slightly over 16 %, or a difference of over 270,000 servers.

Tcl/Tk is the powerful, free, cross-platform scripting language developed by Dr. John Ousterhout, now a Sun Distinguished Engineer. In his own words

"Tcl/Tk lets software developers get the job done ten times faster than with toolkits based on C or C++. It's also a great glue language for making existing applications work together and making them more graphical and Internet-aware."

With a developer community of over 500,000 worldwide, and thousands of commercial applications, Sun has just announced a new business group called SunScript, to support this community with an integrated development environment and to develop a suite of products to link Tcl to the Web and Java.

Karl Lehenbauer, Founder and Chief Technical Officer of NeoSoft, has been part of Tcl/Tk development from the very beginning. Together with Mark Diehkans, they authored Extended Tcl, also known as TclX or NeoSoft Tcl, a powerful set of extensions to the language. Many of the current core Tcl commands originated in Extended Tcl, and were then imported into the core language by Dr. Ousterhout.

NeoSoft Inc., 1770 St. James Place, Suite 500, Houston, TX 77056 USA

19 PHP/FI Server-side html-embedded scripting language for PostgreSQL

WWW Interface Tool `<http://www.vex.net/php>`

Questions e-mail to : rasmus@lerdorf.on.ca

The extract from the home page of PHP/FI is given below:-

PHP/FI is a server-side html-embedded scripting language. It lets you write simple scripts right in your .HTML files much like JavaScript does, except, unlike JavaScript PHP/FI is not browser-dependant. JavaScript is a client-side html-embedded language while PHP/FI is a server-side language. PHP/FI is similar in concept to Netscape's LiveWire Pro product. If you have the money, you run Netscape's Commerce Server and you run one of the supported operating systems, you should probably have a look at LiveWire Pro. If you like free fast-moving software that comes with full source code you will probably like PHP/FI.

19.1 Major Features

Standard CGI, FastCGI and Apache module Support As a standard CGI program, PHP/FI can be installed on any Unix machine running any Unix web server. With support for the new FastCGI standard, PHP/FI can take advantage of the speed improvements gained through this mechanism. As an Apache module, PHP/FI becomes an extremely powerful and lightning fast alternative to CGI programmimg.

- Access Logging With the access logging capabilities of PHP/FI, users can maintain their own hit counting and logging. It does not use the system's central access log files in any way, and it provides real-time access monitoring. The Log Viewer Script provides a quick summary of the accesses to a set of pages owned by an individual user. In addition to that, the package can be configured to generate a footer on every page which shows access information. See the bottom of this page for an example of this.
- Access Control A built-in web-based configuration screen handles access control configuration. It is possible to create rules for all or some web pages owned by a certain person which place various restrictions on who can view these pages and how they will be viewed. Pages can be password protected, completely restricted, logging disabled and more based on the client's domain, browser, e-mail address or even the referring document.
- Postgres Support Postgres is an advanced free RDBMS. PHP/FI supports embedding Postgres95 and PostgreSQL SQL queries directly in .html files.
- RFC-1867 File Upload Support File Upload is a new feature in Netscape 2.0. It lets users upload files to a web server. PHP/FI provides the actual Mime decoding to make this work and also provides the additional framework to do something useful with the uploaded file once it has been received.
- HTTP-based authentication control PHP/FI can be used to create customized HTTP-based authentication mechanisms for the Apache web server.
- Variables, Arrays, Associative Arrays PHP/FI supports typed variables, arrays and even Perl-like associative arrays. These can all be passed from one web page to another using either GET or POST method forms.
- Conditionals, While Loops PHP/FI supports a full-featured C-like scripting language. You can have if/then/elseif/else/endif conditions as well as while loops and switch/case statements to guide the logical flow of how the html page should be displayed.
- Extended Regular Expressions Regular expressions are heavily used for pattern matching, pattern substitutions and general string manipulation. PHP/FI supports all common regular expression operations.
- Raw HTTP Header Control The ability to have web pages send customized raw HTTP headers based on some condition is essential for high-level web site design. A frequent use is to send a Location: URL header to redirect the calling client to some other URL. It can also be used to turn off cacheing or manipulate the last update header of pages.
- On-the-fly GIF image creation PHP/FI has support for Thomas Boutell's GD image library which makes it possible to generate GIF images on the fly.
- ISP "Safe Mode" support PHP/FI supports a unique "Safe Mode" which makes it safe to have multiple users run PHP scripts on the same server.

- It's Free! One final essential feature. The package is completely free. It is licensed under the GPL which allows you to use the software for any purpose, commercial or otherwise. See the GNU Public License document for complete details.

19.2 Credits

* Large parts of this code were developed at and for the University of Toronto. Many thanks to Lee Oattes of the Network Development Department at the university for constant constructive criticism.

* The Postgres95 support code was written by Adam Sussman asussman@vidya.com

* Countless others have helped test and debug the package.

PHP/FI Version 2.0

19.3 Brief History

PHP began life as a simple little cgi wrapper written in Perl. I wrote it in an afternoon during a period between contracts when I needed a quick tool to get an idea of who was reading my online resume. It was never intended to go beyond my own private use. The web server where I had my resume was extremely overloaded and had constant problems forking processes. I rewrote the Perl wrapper in C to get rid of the considerable overhead of having to fork Perl each time my resume was accessed.

Eventually other people on the same web server came across my wrapper and asked if they could use it. Then, as inevitably happens, they started asking for more features. I added more features and finally put together a semi-complete distribution along with documentation, a mailing-list and a FAQ. The name of this first package was Personal Home Page Tools, which later became Personal Home Page Construction Kit.

At the same time I started playing with databases and wrote a tool to easily embed SQL queries into web pages. It was basically another CGI wrapper that parsed SQL queries and made it easy to create forms and tables based on these queries. This tool was named FI (Form Interpreter).

PHP/FI version 2.0 is a complete rewrite of these two packages combined into a single program. It has now evolved to the point where it is a simple programming language embedded inside HTML files. The original acronym, PHP, has stuck. It isn't really appropriate any longer. PHP/FI is used more for entire web sites today than for small Personal Home Page setups. By whatever name, it eliminates the need for numerous small Perl cgi programs by allowing you to place simple scripts directly in your HTML files. This speeds up the overall performance of your web pages since the overhead of forking Perl several times has been eliminated. It also makes it easier to manage large web sites by placing all components of a web page in a single html file. By including support for various databases, it also makes it trivial to develop database enabled web pages. Many people find the embedded nature much easier to deal with than trying to create separate HTML and CGI files.

Throughout this documentation any references to PHP, FI or PHP/FI all refer to the same thing. The difference between PHP and FI is only a conceptual one. Both are built from the same source distribution. When I build the package without any access logging or access restriction support, I call my binary FI. When I build with these options, I call it PHP.

19.4 So, what can I do with PHP/FI?

The first thing you will notice if you run a page through PHP/FI is that it adds a footer with information about the number of times your page has been accessed (if you have compiled access logging into the binary). This is just a very small part of what PHP/FI can do for you. It serves another very important role as a form interpreter cgi, hence the FI part of the name. For example, if you create a form on one of your web pages, you need something to process the information on that form. Even if you just want to pass the information to another web page, you will have to have a cgi program do this for you. PHP/FI makes it extremely easy to take form data and do things with it.

19.5 A simple example

Suppose you have a form:

```
<FORM ACTION="/cgi-bin/php.cgi/~userid/display.html" METHOD=POST>
<INPUT TYPE="text" name="name">
<INPUT TYPE="text" name="age">
<INPUT TYPE="submit">
<FORM>
```

Your display.html file could then contain something like:

< ?echo "Hi $ name, you are $ age years old!

" >

It's that simple! PHP/FI automatically creates a variable for each form input field in your form. You can then use these variables in the ACTION URL file.

The next step once you have figured out how to use variables is to start playing with some logical flow tags in your pages. For example, if you wanted to display different messages based on something the user inputs, you would use if/else logic. In our above example, we can display different things based on the age the user entered by changing our display.html to:

```
<?
    if($age>50);
        echo "Hi $name, you are ancient!<p>";
    elseif($age>30);
        echo "Hi $name, you are very old!<p>";
    else;
        echo "Hi $name.";
    endif;
>
```

PHP/FI provides a very powerful scripting language which will do much more than what the above simple example demonstrates. See the section on the PHP/FI Script Language for more information.

You can also use PHP/FI to configure who is allowed to access your pages. This is done using a built-in configuration screen. With this you could for example specify that only people from certain domains would be allowed to see your pages, or you could create a rule which would password protect certain pages. See the Access Control section for more details.

PHP/FI is also capable of receiving file uploads from any RFC-1867 compliant web browser. This feature lets people upload both text and binary files. With PHP/FI's access control and logical functions, you have full control over who is allowed to upload and what is to be done with the file once it has been uploaded. See the File Upload section for more details.

PHP/FI has support for the Postgres95 database package. It supports embedded SQL queries in your .HTML files. See the section on Postgres95 Support for more information.

PHP/FI also has support for the mysql database package. It supports embedded SQL queries in your .HTML files. See the section on mysql Support for more information.

19.6 CGI Redirection

Apache 1.0.x Notes

A good way to run PHP/FI is by using a cgi redirection module with the Apache server. Please note that you do not need to worry about redirection modules if you are using the Apache module version of PHP/FI. There are two of these redirection modules available. One is developed by Dave Andersen angio@aros.net and it is available at ftp://ftp.aros.net/pub/util/apache/mod_cgi_redirect.c and the other comes bundled with Apache and is called mod_actions.c. The modules are extremely similar. They differ slightly in their usage. Both have been tested and both work with PHP/FI.

One large caveat at the time of this writing (Apr.20/96) is that the current official Apache release (1.0.5) has a severe limitation which prevents cgi redirected requests from having any post-method data associated with them. I have tracked this down and fixed it in my version of Apache, and there is an official patch available in the File Archives on the PHP Home Page.

A second rather large caveat with Apache 1.0.x is that it does not align double types correctly on most architectures. You find find yourself getting strange bus errors from your httpd when using mod_php, either upgrade to Apache 1.1 or edit the alloc.c Apache source file. In this file you will find the following piece of code:

union align { /* Types which are likely to have the longest RELEVANT alignment * restrictions... we don't do much with doubles. */

char *cp; void (*f)(); long l; FILE *fp; };

You will need to add a double to this line and recompile your Apache server. The correct block of code is:

union align { /* Types which are likely to have the longest RELEVANT alignment * restrictions... we don't do much with doubles. */

char *cp; void (*f)(); long l; FILE *fp; double d; };

Check the Apache documentation on how to add a module. Generally you add the module name to a file called Configuration. The line to be added if you want to use the mod_actions module is:

Module action_module mod_actions.o

If you are using the mod_cgi_redirect.c module add this line:

Module cgi_redirect_module mod_cgi_redirect.o

Then compile your httpd and install it. To configure the cgi redirection you need to either create a new mime type in your mime.types file or you can use the AddType command in your srm.conf file to add the mime type. The mime type to be added should be something like this:

application/x-httpd-php phtml

If you are using the mod_actions.c module you need to add the following line to your srm.conf file:

Action application/x-httpd-php /cgi-bin/php.cgi

If you are using mod_cgi_redirect.c you should add this line to srm.conf:

CgiRedirect application/x-httpd-php /cgi-bin/php.cgi

Don't try to use both mod_actions.c and mod_cgi_redirect.c at the same time.

Once you have one of these cgi redirection modules installed and configured correctly, you will be able to specify that you want a file parsed by php/fi simply by making the file's extension .phtml. Furthermore, if you add index.phtml to your DirectoryIndex configuration line in your srm.conf file then the top-level page in a directory will be automatically parsed by php if your index file is called index.phtml.

Netscape HTTPD

You can automatically redirect requests for files with a given extension to be handled by PHP/FI by using the Netscape Server CGI Redirection module. This module is available in the File Archives on the PHP/FI Home Page. The README in the package explicitly explains how to configure it for use with PHP/FI.

NCSA HTTPD

NCSA does not currently support modules, so in order to do cgi redirection with this server you need to modify your server source code. A patch to do this with NCSA 1.5 is available in the PHP/FI file archives.

19.7 Running PHP/FI from the command line

If you build the CGI version of PHP/FI, you can use it from the command line simply typing: php.cgi filename where filename is the file you want to parse. You can also create standalone PHP/FI scripts by making the first line of your script look something like:

```
#!/usr/local/bin/php.cgi -q
```

The "-q" suppresses the printing of the HTTP headers. You can leave off this option if you like.

20 Python Interface for PostgreSQL

PyGres95 is a python interface for the PostgreSQL. It is available from
`<ftp://ftp.via.ecp.fr/pub/python/contrib/Database/PyGres95-1.0b.tar.gz >`

The extract from the home page of PyGres95 is given below:-

PyGres - v1.0b : Postgres95 module for Python

PyGres95, version 1.0b A Python interface for Postgres95 database. Written by Pascal Andre, andre@chimay.via.ecp.fr

Postgres95 is a database system derived from Postgres4.2. It conforms to (most of) ANSI SQL and offer many interesting possibilities (C dynamic linking for functions or type definition, time travel, ...). This package is copyrighted by the Regents of the University of California, and is freely distributable.

Python is a interpretated programming langage. It is object oriented, simple to use (light syntax, simple and straighforward statements), and has many extensions for building GUIs, interfacing with WWW, ... An "intelligent" web browser (HotJava like) is currently under development (november 1995), and this should open programmers many doors. Python is copyrighted by Stichting Mathematisch Centrum, Amsterdam, The Netherlands, and is freely distributable.

PyGres95 is a python module that interfaces Postgres95 database. It embeds Postgres95 query library to allow an easy use of powerful Postgres95 features cooperatively with all the other python modules. It has been developed on a Linux 1.3/ELF system, but have been tested on a Solaris 2.4 platform. Anyway, it should work on any platform where python and postgres95 are available.

20.1 Where to get ... ?

The home sites of the differents packages are:

- Python : `<ftp.python.org:/pub/python>`
- Postgres95 : `<ftp.s2k-ftp.cs.berkeley.edu:/pub/postgres95>`
- PyGres95 : `<ftp.via.ecp.fr:/pub/python/contrib>`

You should anyway try to find some mirror site closer of your site. Refer to the information sources to find these sites. PyGres95 should reside in the contrib directories of Python and Postgres95 sites.

20.2 Information and support

If you need information about these packages please check their web sites:

- Python : `<http://www.python.org/>`
- Postgres95 : `<http://epoch.cs.berkeley.edu:8000/postgres95/index.html>`
- PyGres95 : `<http://www.via.ecp.fr/via/products/pygres.html>`

For support :

- Python : newsgroup comp.lang.python

- Postgres95 : mailing list (see package documentation for information)
- PyGres95 : contact me andre@via.ecp.fr for bug reports, ideas, remarks

I will try to answer as long as my free time allow me to do that.

21 Gateway between PostgreSQL and the WWW - WDB-P95

WDB-P95 - A Web interface to Postgres95 Databases. It is at
`<http://www.eol.ists.ca/ dunlop/wdb-p95/>`

The extract from the home page of WDB-P95 is given below:-

Version 1.4b2 beta - Created by J. Douglas Dunlop

About wdb-p95

This is a modified version of wdb-1.3a2 which provides a gateway to a the WWW for Postgres95. This version also requires a Browser that is capable of handling HTML Tables for the tabular output. This is not required by the original wdb and can be fairly easily reverted. (I only wanted tables because the < pre > stuff just didn't agree with me!)

You can try out my CASI Tape and Image Query. You can have a peek at the Form Definition File (FDF) that I used to create the CASI Tape and Image Query too, which includes a JOIN of 2 tables.

This release contains all files necessary to install and run WDB-P95 as an interface to your Postgres95 databases. To port this system to other database should be relatively easy - provided that it supports standard SQL and has a Perl interface.

21.1 Does the Postgres95 server,pgperl, and httpd have to be on the same host?

No - the Postgres95 server does not have to be on the same host. As WDB-P95 is called by the http daemon, they have to be on the same host. - And as WDB-P95 was written to use Pg.pm - pgperl has to be on the same host too. Pgperl was written using the libpq library, so it will be able to access any Postgres95 server anywhere in the net, just like any other Postgres95 client. As illustrated below

{WWW Client (Netscape)} =¿ {HTTP Server (NCSA's http) + WDB-P95 + pgperl + libpq}=¿ {Postgres95 server}

Curly brackets {} represent machines.

Each machine can be of a different type : NT, SUN, HP, ... but you need the libpq interface library for the machine type where you plan to use WDB-P95, as you need it to compile pgperl. (The system was designed to use HTML tables so a recent WWW client is best)

21.2 New Version

New versions of the software and the above pages are always available from the WDB-P95 Home page. `<http://www.eol.ists.ca/ dunlop/wdb-p95/>`

For questions or to join Mailing lists contact dunlop@eol.ists.ca

22 "C" language Interface for PostgreSQL

Included in distribution and is called 'libpq'. Similar to Oracle OCI, Sybase DB-lib, Informix CLI libraries.

23 "C++" language Interface for PostgreSQL

Included in distribution and is called 'libpq++'.

24 ESQL/C for PostgreSQL

Embedded C Pre-compiler for PostgreSQL ESQL/C like Oracle Pro*C, Informix ESQL/C:

- `<ftp://ftp.lysator.liu.se/pub/linus>`
- Email : linus@epact.se

The PostgreSQL ESQL/C is an SQL application-programming interface (API) enables the C programmer to create custom applications with database-management capabilities. The PostgreSQL ESQL/C allows you to use a third-generation language with which you are familiar and still take advantage of the Structured Query Language (SQL).

ESQL/C consists of the following pieces of software:

- The ESQL/C libraries of C functions provide access to the database server.
- The ESQL/C header files provide definitions for the data structures, constants, and macros useful to the ESQL/C program.
- The ESQL/C preprocessor, is a source-code preprocessor that converts a C file containing SQL statements into an executable file.

25 Japanese Kanji Code for PostgreSQL

Very useful for Japanese people. It is at the following site `<ftp://ftp.sra.co.jp/pub/cmd/postgres/>`

26 PostgreSQL Port to Windows 95/Windows NT

Port to Windows 95/Windows NT is underway. Porting is being done using gcc, gmake for Win NT/95. To compile source code on win32 gnu-win32 program is used. GNU gcc is available for win32. Check this site -

- `<http://www.cygnus.com/misc/gnu-win32>`

At this site and get the file cdk.exe (self-extractor file for gnu-win32)

Porting can also be done using the following "Unix-Emulator on NT" tool from

- `<http://www.softway.com>`

27 Mailing Lists

See the Mailing Lists Item on the main web page at : `<http://www.postgresql.org/>`

- Email questions to: pgsql-questions@postgresql.org
- Developers pgsql-hackers@postgresql.org
- Port specific questions pgsql-ports@postgresql.org
- Documentation questions pgsql-docs@postgresql.org

You will get the answers/replies back by e-mail in less than a day!!

You can also subscribe to mailing lists. To subscribe or unsubscribe from the list, send mail to

- pgsql-questions-request@postgresql.org
- pgsql-hackers-request@postgresql.org
- pgsql-ports-request@postgresql.org
- pgsql-docs-request@postgresql.org

The body of the message should contain the single line

subscribe

(or)

unsubscribe

Also mailing lists are archived in html format at the following location -

- `<ftp://ftp.postgresql.org>` directory is /pub/majordomo

28 Documentations and Books

Included in the distribution is the

- 'User Guide' for PostgreSQL,
- 'Implementation Guide' detailing database internals of PostgreSQL.
- Online manuals.
- Online manuals in HTML formats.
- Also manuals in Postscript format for printing hard copies.

Reference docs: Useful reference textbooks :

- "Understanding the New SQL: A Complete Guide" - by Jim Melton and Alan R.Simon Morgan Kaufman Publisher is one of best SQL books.
- "A Guide to THE SQL STANDARD" - by C.J.Date Addison-Wesley Publishing company is also a good book

Hundreds of other titles on SQL are available! Check out a bookstore.

29 Technical support for PostgreSQL

- You can e-mail your technical questions or problems you face to: pgsql-questions@postgresql.org

and you will get back e-mail answer in less than a day.

You can also purchase technical support from the following companies/corporations -

- `<http://www.redhat.com>`
- `<http://www.caldera.com>`
- `<http://www.wgs.com>`
- `<http://www.yggdrasil.com>`

30 Economic and Business Aspects

Commercial databases pay federal state taxes, sales tax, employment taxes, social security taxes, health care for employees, bunch of benefits for employees, marketing and advertisement costs. All these costs do not go directly for the development of the database. When you buy a commercial database, about 60% of amount goes to taxes+benefits, 20% for marketing, 10% for profit margin and the balance 10% for actual database R&D costs. Hence the real worth of the database is much less.

Also commercial databases have to pay for buildings/real-estates and purchase Unix machines, install and maintain them. All of these costs are passed onto customers.

PostgreSQL has the advantage over commercial databases as there is no direct tax since it is made on the internet. A very vast group of companies contribute to the development of the PostgreSQL. For example, if there are one million companies in U.S.A and each contribute about $ 10 (worth of software to PostgreSQL) than each and every company will get ten million dollars!! This is the magic of software development on internet.

Currently, PostgreSQL source code is about 2,00,000 lines of "C", "C++" code. If cost of each line of "C" code is $ 10 than the total cost of PostgreSQL as of today is $ 2,000,000 (two million dollars!!).

Many companies already develop in-house vast amount of "C", "C++" code. Hence by taking in the source code of PostgreSQL and collaborating with other companies on internet will greatly benefit the company saving time and efforts.

31 Conclusion

After researching all the available databases which are free and source code is available, it was found that ONLY PostgreSQL is the MOST mature, most widely used and robust RDBMS SQL free database in the world. PostgreSQL is very appealing as lot of work had already been done and it has ODBC, JDBC drivers using these it is possible to write applications independent of the databases. The applications written in PostgreSQL using ODBC, JDBC drivers are easily portable to other databases like Oracle, Sybase and Informix. And applications written for Oracle, Sybase and Informix using ODBC, JDBC drivers are easily portable to PostgreSQL database.

32 FAQ - Questions on PostgreSQL

32.1 Frequently Asked Questions (FAQ) for PostgreSQL

Last updated: Wed Jun 11 10:44:40 EDT 1997 Version: 6.1.1

Current maintainer: Bruce Momjian (maillist@candle.pha.pa.us)

The most recent version of this document can be viewed at the postgreSQL Web site, `<http://www.postgreSQL.org>`

Linux-specific questions are answered in `<http://www.postgreSQL.org/docs/FAQ-Linux.phtml>`

Irix-specific questions are answered in `<http://www.postgreSQL.org/docs/FAQ-Irix.phtml>`

Changes in this version (* = modified, + = new):

* 3.42) What is Genetic Query Optimization? * 3.43) I am running Solaris and my dates display wrong. Why?

32.2 What is PostgreSQL?

PostgreSQL is an enhancement of the POSTGRES database management system, a next-generation DBMS research prototype. While PostgreSQL retains the powerful data model and rich data types of POSTGRES, it replaces the PostQuel query language with an extended subset of SQL. PostgreSQL is free and the complete source is available.

PostgreSQL development is being performed by a team of Internet developers who all subscribe to the PostgreSQL development mailing list. The current coordinator is Marc G. Fournier (scrappy@postgreSQL.org). (See below on how to join). This team is now responsible for all current and future development of PostgreSQL.

The authors of PostgreSQL 1.01 were Andrew Yu and Jolly Chen. Many others have contributed to the porting, testing, debugging and enhancement of the code. The original Postgres code, from which PostgreSQL is derived, was the effort of many graduate students, undergraduate students, and staff programmers working under the direction of Professor Michael Stonebraker at the University of California, Berkeley.

The original name of the software at Berkeley was Postgres. When SQL functionality was added in 1995, its name was changed to Postgres95. The name was changed at the end of 1996 to PostgreSQL.

32.3 What does PostgreSQL run on?

The authors have compiled and tested PostgreSQL on the following platforms(some of these compiles require gcc 2.7.0):

* aix - IBM on AIX 3.2.5

* alpha - DEC Alpha AXP on OSF/1 2.0

* BSD44_derived - OSs derived from 4.4-lite BSD (NetBSD, FreeBSD)

* bsdi - BSD/OS 2.0, 2.01, 2.1, 3.0

* dgux - DG/UX 5.4R3.10

* hpux - HP PA-RISC on HP-UX 9.0

* i386_solaris - i386 Solaris

* irix5 - SGI MIPS on IRIX 5.3

* linux - Intel x86 on Linux 1.2 and Linux ELF (For non-ELF Linux, see LINUX_ELF below).

* sparc_solaris - SUN SPARC on Solaris 2.4

* sunos4 - SUN SPARC on SunOS 4.1.3

* svr4 - Intel x86 on Intel SVR4 * ultrix4 - DEC MIPS on Ultrix 4.4

The following platforms have known problems/bugs:

* nextstep - Motorola MC68K or Intel x86 on NeXTSTEP 3.2

32.4 Where can I get PostgreSQL?

The primary anonymous ftp site for PostgreSQL is:

* <ftp://ftp.postgreSQL.org/pub>

A mirror site exists at:

* <ftp://postgres95.vnet.net/pub/postgres95> * <ftp://ftp.luga.or.at/pub/postgres95> * <ftp://cal011111.student.utwente.nl/pub/postgres95> * <ftp://ftp.uni-trier.de/pub/database/rdbms/postgres/postgres95> * <ftp://rocker.sch.bme.hu>

32.5 What's the copyright on PostgreSQL?

PostgreSQL is subject to the following COPYRIGHT.

PostgreSQL Data Base Management System

Copyright (c) 1994-6 Regents of the University of California

Permission to use, copy, modify, and distribute this software and its documentation for any purpose, without fee, and without a written agreement is hereby granted, provided that the above copyright notice and this paragraph and the following two paragraphs appear in all copies.

IN NO EVENT SHALL THE UNIVERSITY OF CALIFORNIA BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING LOST PROFITS, ARISING OUT OF THE USE OF THIS SOFTWARE AND ITS DOCUMENTATION, EVEN IF THE UNIVERSITY OF CALIFORNIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE UNIVERSITY OF CALIFORNIA SPECIFICALLY DISCLAIMS ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE SOFTWARE PROVIDED HEREUNDER IS ON AN "AS IS" BASIS, AND THE UNIVERSITY OF CALIFORNIA HAS NO OBLIGATIONS TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

32.6 Support for PostgreSQL

There is no official support for PostgreSQL from the original maintainers or from University of California, Berkeley. It is maintained through volunteer effort only.

The main mailing list is: questions@postgreSQL.org. It is available for discussion o f matters pertaining to PostgreSQL, including but not limited to bug reports and fixes. For info on how to subscribe, send a mail with the lines in the body (not the subject line)

subscribe end

to questions-request@postgreSQL.org.

There is also a digest list available. To subscribe to this list, send email to:
questions-digest-request@postgreSQL.org with a BODY of:

subscribe end

Digests are sent out to members of this list whenever the main list has received around 30k of messages.

There is a bugs mailing list available. To subscribe to this list, send email to bugs-request@postgreSQL.org with a BODY of:

There is also a developers discussion mailing list available. To subscribe to this list, send email to hackers-request@postgreSQL.org with a BODY of:

subscribe end

Additional information about PostgreSQL can be found via the PostgreSQL WWW home page at:

`<http://www.postgreSQL.org>`

32.7 Latest release of PostgreSQL

The latest release of PostgreSQL is version 6.0, which was released on January 31, 1997. 6.1 is scheduled for release soon. For information about what is new in 6.1, see our TODO list on our WWW page.

We expect a 7.0 release in several months that will remove time-travel and reduce by 50 % the size of on-disk system columns maintained for each row in a table. This release will also require a dump and restore.

32.8 Is there a commercial version of PostgreSQL?

Illustra Information Technology (a wholly owned subsidiary of Informix Software, Inc.) sells an object-relational DBMS called Illustra that was originally based on postgres. Illustra has cosmetic similarities to PostgreSQL but has more features, is more robust, performs better, and offers real documentation and support. On the flip side, it costs money. For more information, contact sales@illustra.com

32.9 What documentation is available for PostgreSQL?

A user manual, manual pages, and some small test examples are included in the distribution. The sql and built-in manual pages are particularly important.

The www page contains pointers to an implementation guide and five papers written about postgres design concepts and features.

32.10 What version of SQL does PostgreSQL use?

PostgreSQL supports a subset of SQL-92. It has most of the important constructs but lacks some of the functionality. The most visible differences are:

* no support for nested subqueries * no HAVING clause under a GROUP BY

On the other hand, you get to create user-defined types, functions, inheritance etc. If you're willing to help with PostgreSQL coding, eventually we can also add the missing features listed above.

32.11 Does PostgreSQL work with databases from earlier versions of postgres?

PostgreSQL v1.09 is compatible with databases created with v1.01. Those upgrading from 1.0 should read the directions in the MIGRATION_1.0_TO_1.02 directory.

Upgrading to 6.0 requires a dump and restore from previous releases.

Upgrading to 6.1 requires a dump and restore from previous releases.

Those ugrading from versions earlier than 1.09 must upgrade to 1.09 first without a dump/reload, then dump the data from 1.09, and then load it into 6.0 or 6.1.

32.12 How many people use PostgreSQL?

Since we don't have any licensing or registration scheme, it's impossible to tell. We do know hundreds copies of PostgreSQL v1.* have been downloaded, and that there many hundreds of subscribers to the mailing lists.

———————————————————————————————

33 FAQ - Installation Questions

33.1 Initdb doesn't run

* check to see that you have the proper paths set * check that the 'postgres' user owns all the right files * ensure that there are files in $ PGDATA/files, and that they are non-empty. If they aren't, then "gmake install" failed for some reason

33.2 When I start up the postmaster, I get "FindBackend- could not find a backend to execute..." "postmaster- could not find backend to execute..."

You probably do not have the right path set up. The 'postgres' executable needs to be in your path.

33.3 The system seems to be confused about commas, decimal points, and date formats.

Check your locale configuration. PostgreSQL uses the locale settings of the user that ran the postmaster process. Set those accordingly for your operating environment.

33.4 How do I install PostgreSQL somewhere other than /usr/local/pgsql?

You need to edit Makefile.global and change POSTGRESDIR accordingly, or create a Makefile.custom and define POSTGRESDIR there.

33.5 When I run postmaster, I get a Bad System Call core dumped message.

It could be a variety of problems, but first check to see that you have system V extensions installed on your kernel. PostgreSQL requires kernel support for shared memory.

33.6 When I try to start the postmaster, I get IpcMemoryCreate errors.

You either do not have shared memory configured properly in kernel or you need to enlarge the shared memory available in the kernel. The exact amount you need depends on your architecture and how many buffers you configure postmaster to run with. For most systems, with default buffer sizes, you need a minimum of ~ 760K.

33.7 I have changed a source file, but a recompile does not see the change?

The Makefiles do not have the proper dependencies for include files. You have to do a 'make clean' and then another 'make'.

---

34 FAQ - PostgreSQL Features

34.1 How do I specify a KEY or other constraints on a column?

Column constraints are not supported in PostgreSQL. As a consequence, the system does not check for duplicates.

Under 6.0, create a unique index on the column. Attempts to create duplicate of that column will report an error.

34.2 Does PostgreSQL support nested subqueries?

Subqueries are not implemented, but they can be simulated using sql functions.

34.3 How do I define a unique indices?

PostgreSQL 6.0 supports unique indices.

34.4 I've having a lot of problems using rules.

Currently, the rule system in PostgreSQL is mostly broken. It works enough to support the view mechanism, but that's about it. Use PostgreSQL rules at your own peril.

34.5 I can't seem to write into the middle of large objects reliably.

The Inversion large object system in PostgreSQL is also mostly broken. It works well enough for storing large wads of data and reading them back out, but the implementation has some underlying problems. Use PostgreSQL large objects at your own peril.

34.6 Does PostgreSQL have a graphical user interface? A report generator? A embedded query language interface?

No. No. No. Not in the official distribution at least. Some users have reported some success at using 'pgbrowse' and 'onyx' as frontends to PostgreSQL. Several contributions are working on tk based frontend tools. Ask on the mailing list.

34.7 How can I write client applications to PostgreSQL?

PostgreSQL supports a C-callable library interface called libpq as well as a Tcl-based library interface called libtcl.

Others have contributed a perl interface and a WWW gateway to PostgreSQL. See the PostgreSQL home pages for more details.

34.8 How do I prevent other hosts from accessing my PostgreSQL backend?

Use host-based authentication by modifying the file $ PGDATA/pg_hba accordingly.

34.9 How do I set up a pg_group?

Currently, there is no easy interface to set up user groups. You have to explicitly insert/update the pg_group table. For example:

```
jolly=> insert into pg_group (groname, grosysid, grolist)
jolly=>     values ('posthackers', '1234', '5443, 8261');
INSERT 548224
jolly=> grant insert on foo to group posthackers;
CHANGE
jolly=>
```

The fields in pg_group are:

* groname: the group name. This a char16 and should be purely alphanumeric. Do not include underscores or other punctuation. * grosysid: the group id. This is an int4. This should be unique for each group. * grolist: the list of pg_user id's that belong in the group.

```
This is an int4[].
```

34.10 What is the exact difference between binary cursors and normal cursors?

Normal cursors return data back in ASCII format. Since data is stored natively in binary format, the system must do a conversion to produce the ASCII format. In addition, ASCII formats are often large in size than binary format. Once the attributes come back in ASCII, often the client application then has to convert it to a binary format to manipulate it anyway.

Binary cursors give you back the data in the native binary representation. Thus, binary cursors will tend to be a little faster since there's less overhead of conversion.

However, ASCII is architectural neutral whereas binary representation can differ between different machine architecture. Thus, if your client machine uses a different representation than you server machine, getting back attributes in binary format is probably not what you want. Also, if your main purpose is displaying the data in ASCII, then getting it back in ASCII will save you some effort on the client side.

34.11 Why doesn't the != operator work?

SQL specifies $< >$ as the inequality operator, and that is what we have defined for the built-in types.

In 6.0, != is equivalent to $< >$.

34.12 What is a R-tree index and what is it used for?

An r-tree index is used for indexing spatial data. A hash index can't handle range searches. A B-tree index only handles range searches in a single dimension. R-tree's can handle multi-dimensional data. For example, if a R-tree index can be built on an attribute of type 'point', the system can more efficient answer queries like select all points within a bounding rectangle.

The canonical paper that describes the original R-Tree design is:

Guttman, A. "R-Trees: A Dynamic Index Structure for Spatial Searching." Proc of the 1984 ACM SIGMOD Int'l Conf on Mgmt of Data, 45-57.

You can also find this paper in Stonebraker's "Readings in Database Systems"

34.13 What is the maximum size for a tuple?

Tuples are limited to 8K bytes. Taking into account system attributes and other overhead, one should stay well shy of 8,000 bytes to be on the safe side. To use attributes larger than 8K, try using the large objects interface.

Tuples do not cross 8k boundaries so a 5k tuple will require 8k of storage.

34.14 I defined indices but my queries don't seem to make use of them. Why?

PostgreSQL does not automatically maintain statistics. One has to make an explicit 'vacuum' call to update the statistics. After statistics are updated, the optimizer has a better shot at using indices. Note that the optimizer is limited and does not use indices in some circumstances (such as OR clauses).

If the system still does not see the index, it is probably because you have created an index on a field with the improper *_ops type. For example, you have created a CHAR(4) field, but have specified a char_ops index type_class.

See the create_index manual page for information on what type classes are available. It must match the field type.

Postgres does not warn the user when the improper index is created.

Indexes not used for ORDER BY operations.

34.15 Are there ODBC drivers for PostgreSQL?

There are two ODBC drivers available, PostODBC and OpenLink ODBC.

For all people being interested in PostODBC, there are now two mailing lists devoted to the discussion of PostODBC. The mailing lists are:

* postodbc-users@listserv.direct.net * postodbc-developers@listserv.direct.net

these lists are ordinary majordomo mailing lists. You can subscribe by sending a mail to:

* majordomo@listserv.direct.net

OpenLink ODBC is currently in beta under Linux. You can get it from
`<http://www.openlinksw.com/postgres.html>` It works with our standard ODBC client software so you'll have Postgres ODBC available on every client platform we support (Win, Mac, Unix, VMS).

We will probably be selling this product to people who need commercial-quality support, but a freeware version will always be available. Questions to postgres95@openlink.co.uk.

34.16 How do I use postgres for multi-dimensional indexing (¿ 2 dimensions)?

Builtin R-Trees can handle polygons and boxes. In theory, R-trees can be extended to handle higher number of dimensions. In practice, extending R-trees require a bit of work and we don't currently have any documentation on how to do it.

34.17 How do I do regular expression searches? case-insensitive regexp searching?

```
PostgreSQL supports the SQL LIKE syntax as well as more general regular
expression searching with the ~ operator. The !~ is the negated regexp
operator. ~* and !~* are the case-insensitive regular expression operators.
```

34.18 I can't access the database as the 'root' user.

You should not create database users with user id 0(root). They will be unable to access the database. This is a security precaution because of the ability of any user to dynamically link object modules into the database engine.

34.19 I experienced a server crash during a vacuum. How do I remove the lock file?

If the server crashes during a vacuum command, chances are it will leave a lock file hanging around. Attempts to re-run the vacuum command result in

WARN:can't create lock file – another vacuum cleaner running?

If you are sure that no vacuum is actually running, you can remove the file called "pg_vlock" in your database directory (which is $ PGDATA/base/< dbName >)

34.20 What is the difference between the various character types?

```
Type            Internal Name   Notes

--------------------------------------------------

CHAR            char            1 character

CHAR2           char2           2 characters

CHAR4           char4           4 characters   optimized for a fixed length

CHAR8           char8           8 characters

CHAR16          char16          16 characters

CHAR(#)         bpchar          blank padded to the specified fixed length

VARCHAR(#)      varchar         size specifies maximum length, no padding

TEXT            text            length limited only by maximum tuple length
```

```
BYTEA           bytea           variable-length array of bytes

Remember, you need to use the internal name when creating indexes on these
fields or when doing other internal operations.

The last four types above are "varlena" types (i.e. the first four bytes is
the length, followed by the data). CHAR(#) and VARCHAR(#) allocate the
maximum number of bytes no matter how much data is stored in the field.
TEXT and BYTEA are the only character types that have variable length on
the disk.
```

34.21 In a query, how do I detect if a field is NULL?

PostgreSQL has two builtin keywords, "isnull" and "notnull" (note no spaces). Version 1.05 and later and 6.* understand IS NULL and IS NOT NULL.

34.22 How do I see how the query optimizer is evaluating my query?

Place the word 'EXPLAIN' at the beginning of the query, for example:

EXPLAIN SELECT * FROM table1 WHERE age = 23;

34.23 How do I create a serial field?

Postgres does not allow the user to specifiy a user column as type SERIAL. Instead, you can use each row's oid field as a unique value. However, if you need to dump and reload the database, you need to be using postgres version 1.07 or later or 6.* with pgdump's -o option or COPY's WITH OIDS option to preserver the oids.

Another valid way of doing this is to create a function:

create table my_oids (f1 int4);

insert into my_oids values (1);

create function new_oid () returns int4 as 'update my_oids set f1 = f1 + 1; select f1 from my_oids; ' language 'sql';

then:

create table my_stuff (my_key int4, value text);

insert into my_stuff values (new_oid(), 'hello');

However, keep in mind there is a race condition here where one server could do the update, then another one do an update, and they both could select the same new id. This statement should be performed within a transaction.

Sequences are implemented in 6.1

34.24 How do I create a multi-column index?

In 6.0, you can not directly create a multi-column index using create index. You need to define a function which acts on the multiple columns, then use create index with that function.

In 6.1, this feature is available.

34.25 What are the temp_XXX files in my database directory?

They are temp_ files generated by the query executor. For example, if a sort needs to be done to satisfy an ORDER BY, some temp files are generated as a result of the sort.

If you have no transactions or sorts running at the time, it is safe to delete the temp_ files.

34.26 Why are my table files not getting any smaller after a delete?

If you run vacuum in pre-6.0, unused rows will be marked for reuse, but the file blocks are not released.

In 6.0, vacuum properly shrinks tables.

34.27 Why can't I connect to my database from another machine?

The default configuration allows only connections from tcp/ip host localhost. You need to add a host entry to the file pgsql/data/pg_hba.

34.28 I get the error 'default index class unsupported' when creating an index. How do I do it?

You probably used:

create index idx1 on person using btree (name);

PostgreSQL indexes are extensible, and therefore in pre-6.0, you must specify a class_type when creating an index. Read the manual page for create index (called create_index).

Version 6.0, if you do not specify a class_type, it defaults to the proper type for the column.

34.29 Why does creating an index crash the backend server?

You have probably defined an incorrect *_ops type class for the field you are indexing.

34.30 How do I find out what indexes or operations are defined in the database?

Run the file pgsql/src/tutorial/syscat.source. It illustrates many of the 'select's needed to get information out of the database system tables.

34.31 Why do statements require an extra character at the end? Why does 'createuser' return 'unexpected last match in input()'? Why does pg_dump fail?

You have compile postgres with flex version 2.5.3. There is bug in this version of flex. Use flex version 2.5.2 or flex 2.5.4 instead. There is a doc/README.flex file which will properly patch the flex 2.5.3 source code.

34.32 All my servers crash under concurrent table access. Why?

This problem can be caused by a kernel that is not configured to support semaphores.

34.33 What tools are available for hooking postgres to Web pages?

For web integration, PHP/FI is an excellent interface. The URL for that is `<http://www.vex.net/php/>`

PHP is great for simple stuff, but for more complex stuff, some still use the perl interface and CGI.pm.

An example of using WWW with C to talk to Postgres is can be tried at:

* `<http://www.postgreSQL.org/ mlc>`

An WWW gatway based on WDB using perl can be downloaded from:

* `<http://www.eol.ists.ca/ dunlop/wdb -p95>`

34.34 What is the time-warp feature and how does it relate to vacuum?

PostgreSQL handles data changes differently than most database systems. When a row is changed in a table, the original row is marked with the time it was changed, and a new row is created with the current data. By default, only current rows are used in a table. If you specify a date/time after the table name in a FROM clause, you can access the data that was current at that time, i.e.

SELECT * FROM employees $'July24, 199609:00:00'$

displays employee rows in the table at the specified time. You can specify intervals like $date, date$, $date$,, , $date$, or ,. This last option accesses all rows that ever existed.

INSERTed rows get a timestamp too, so rows that were not in the table at the desired time will not appear.

Vacuum removes rows that are no longer current. This time-warp feature is used by the engine for rollback and crash recovery. Expiration times can be set with purge.

In 6.0, once a table is vacuumed, the creation time of a row may be incorrect, causing time-traval to fail.

The time-travel feature will be removed in 7.0.

34.35 How do I tune the database engine for better performance?

There are two things that can be done. You can use Openlink's option to disable fsync() by starting the postmaster with a '-o -F' option. This will prevent fsync()'s from flushing to disk after every transaction.

You can also use the postmaster -B option to increase the number of shared memory buffers shared among the backend processes. If you make this parameter too high, the process will not start or crash unexpectedly. Each buffer is 8K and the defualt is 64 buffers.

34.36 What debugging features are available in PostgreSQL?

PostgreSQL has several features that report status information that can be valuable for debugging purposes.

First, by compiling with DEBUG defined, many assert()'s monitor the progress of the backend and halt the program when something unexpected occurs.

Both postmaster and postgres have several debug options available. First, whenever you start the postmaster, make sure you send the standard output and error to a log file, like:

```
cd /usr/local/pgsql
./bin/postmaster >server.log 2>&1 &
```

This will put a server.log file in the top-level PostgreSQL directory. This file can contain useful information about problems or errors encountered by the server. Postmaster has a -d option that allows even more detailed information to be reported. The -d option takes a number 1-3 that specifies the debug level. The query plans in a verbose debug file can be formatted using the 'indent' program. (You may need to remove the '====' lines in 1.* releases.) Be warned that a debug level greater than one generates large log files in 1.* releases.

You can actuall run the postgres backend from the command line, and type your SQL statement directly. This is recommended ONLY for debugging purposes. Note that a newline terminates the query, not a semicolon. If you have compiled with debugging symbols, you can perhaps use a debugger to see what is happening. Because the backend was not started from the postmaster, it is not running in an identical environment and locking/backend interaction problems may not be duplicated. Some operating system can attach to a running backend directly to diagnose problems.

The postgres program has a -s, -A, -t options that can be very usefull for debugging and performance measurements.

The EXPLAIN command (see this FAQ) allows you to see how PostgreSQL is iterpreting your query.

34.37 What is an oid? What is a tid?

Oids are Postgres's answer to unique row ids or serial columns. Every row that is created in Postgres gets a unique oid. All oids generated by initdb are less than 16384 (from backend/access/transam.h). All post-initdb (user-created) oids are equal or greater that this. All these oids are unique not only within a table, or database, but unique within the entire postgres installation.

Postgres uses oids in its internal system tables to link rows in separate tables. These oids can be used to identify specific user rows and used in joins. It is recommended you use column type oid to store oid values. See the sql(l) manual page to see the other internal columns.

Tids are used to indentify specific physical rows with block and offset values. Tids change after rows are modified or reloaded. They are used by index entries to point to physical rows. They can not be accessed through sql.

34.38 What is the meaning of some of the terms used in Postgres?

Some of the source code and older documentation use terms that have more common usage. Here are some:

* row, record, tuple * attribute, field, column * table, class * retrieve, select * replace, update * append, insert * oid, serial value * portal, cursor * range variable, table name, table alias

Please let me know if you think of any more.

34.39 What is Genetic Query Optimization?

The GEQO module in PostgreSQL is intended to solve the query optimization problem of joining many tables by means of a Genetic Algorithm (GA). It allows the handling of large join queries through non-exhaustive search.

For further information see README.GEQO utesch@aut.tu-freiberg.de.

34.40 I am running Solaris and my dates display wrong. Why?

There was a bug in 6.0 that caused this problem under Solaris with -O2 optimization. Upgrade to 6.1.

34.41 How do I enable more than 32 concurrent backends?

Edit include/storage/sinvaladt.h, and change the value of MaxBackendId. In the future, we plan to make this a configurable prameter.

---

35 FAQ - Extending PostgreSQL

35.1 I wrote a user-defined function and when I run it in psql, it dumps core.

The problem could be a number of things. Try testing your user-defined function in a stand alone test program first. Also, make sure you are not sending elog NOTICES when the front-end is expecting data, such as during a type_in() or type_out() functions

35.2 I get messages of the type NOTICE-PortalHeapMemoryFree-0x402251d0 not in alloc set!

You are pfree'ing something that was not palloc'ed. When writing user-defined functions, do not include the file "libpq-fe.h". Doing so will cause your palloc to be a malloc instead of a free. Then, when the backend pfrees the storage, you get the notice message.

35.3 I've written some nifty new types and functions for PostgreSQL.

Please share them with other PostgreSQL users. Send your extensions to mailing list, and they will eventually end up in the contrib/ subdirectory.

35.4 How do I write a C function to return a tuple?

This requires extreme wizardry, so extreme that the authors have not ever tried it, though in principle it can be done. The short answer is ... you can't. This capability is forthcoming in the future.

—————————————————————————————

36 FAQ - Bugs

36.1 How do I make a bug report?

Check the current FAQ at `<http://www.postgreSQL.org>`

Also check out our ftp site `<ftp://ftp.postgreSQL.org/pub>` to see if there is a more recent PostgreSQL version.

You can also fill out the "bug-template" file and send it to:

* bugs@postgreSQL.org

This is the address of the developers mailing list.

37 FAQ Linux-PostgreSQL - Compiling PostgreSQL

==

Frequently Asked Questions (FAQ) for PostgresSQL V6.0 Linux Specific TO BE READ IN CONJUNCTION WITH THE NORMAL FAQ

==

last updated: Wed Jan 29 20:16:00 GMT 1997

current maintainer: Andrew C.R. Martin (martin@biochem.ucl.ac.uk)

original author: Andrew C.R. Martin (martin@biochem.ucl.ac.uk)

37.1 What changes do I need to make to src/Makefile.global or src/Makefile.custom?

These changes are most easily made by running the customize shell script in the src directory which will write a Makefile.custom for you.

If you do it by hand, you *must* set the following variable: PORTNAME= linux

You will also need to change the following to match your own installation: POSTGRESDIR

If you switch on the USE_TCL option, you will need to set these: TCL_INCDIR= TCL_LIBDIR= TCL_LIB= TK_INCDIR= TK_LIBDIR= TK_LIB= X11_INCDIR= X11_LIBDIR= X11_LIB=

On my Slackware3.0 system, these are: TCL_INCDIR= /usr/include/tcl TCL_LIBDIR= /usr/lib TCL_LIB= -ltcl TK_INCDIR= /usr/include/tcl TK_LIBDIR= /usr/lib TK_LIB= -ltk X11_INCDIR= /usr/include/X11 X11_LIBDIR= /usr/X386/lib X11_LIB= -lX11

You may also make any other changes you need as documented in the INSTALL file and in Makefile.global

37.2 Why do I get problems with missing libreadline?

Linux systems generally don't come with the GNU readline library installed. Either ensure you do not activate the readline options in src/Makefile.global or src/Makefile.custom or install the GNU readline library.

Note that Debian Linux (like FreeBSD) does come with readline installed.

37.3 *REDHAT* Why do I get problems with missing libdl and dlfcn.h?

The libdl library is used for dynamic linking of user-supplied functions at run-time. For some reason this library was missed out from the Redhat distribution. It seems that the latest Redhat 4.0 (Colgate) fixes this.

RedHat now have a new ld.so RPM package on their FTP site. Simply grab:

`<ftp://ftp.redhat.com/pub/redhat/devel/i386/RedHat/RPMS/ld.so-1.7.14-4.i386.rpm>`

Install the RPM file in the usual way and off you go!

There has been one report of a corrupted system resulting from programs accessing these libraries while updating them (not altogether surprising). Consequently it is a good idea to reboot the system before installing the new libraries and to have as little running as possible during this upgrade. Going into single-user mode is probably a good idea!

If you want to do it the hard way, you can obtain the library and the header file from:

`<ftp://tsx-11.mit.edu/pub/linux/packages/GCC/ld.so-1.7.14.tar.gz>`

Alternatively, you may find precompiled binaries in distributions/debian/buzz/binary-i386/base/ld.so-1.7.14-4.deb on the same site, or follow the instructions given for question 1.2 for correcting the same error with early releases of Slackware 3.1. Don't use this method unless you know what you are doing!

37.4 *SLACKWARE*3.1 Why do I get problems with missing libdl and dlfcn.h?

See the answer to question 1.3. Slackware up to version 3.0 was supplied with this library and include file and they seem to be back in again in the latest versions of 3.1, but the early 3.1 releases (before 9th September 1996) had them missing and many CD-ROM versions will have been pressed from the first 3.1 releases.

There has been one report of a corrupted system resulting from programs accessing these libraries while updating them (not altogether surprising). Consequently it is a good idea to reboot the system before installing the new libraries and to have as little running as possible during this upgrade. Going into single-user mode is probably a good idea!

The easiest fix is to obtain the file ldso.tgz from the a4 disk of a more recent Slackware 3.1 distribution and unpack this file from the root (/) directory, then do

sh install/doinst.sh

to complete the installation. Follow this with

ldconfig

If you want to install manually, you must first install the file dlfcn.h in /usr/include.

Second, install the file libdl.so.1.7.14 (or whatever the latest release is) in /lib, then do:

cd /lib ln -sf libdl.so.1.7.14 libdl.so.1 ln -sf libdl.so.1 libdl.so

On some systems (depending on your GCC configuration) it may be necessary to do:

cd /usr/lib ln -sf /lib/libdl.so .

Finally

ldconfig

37.5 My compile of the backend dies complaining about the include file dlfcn.h missing

See the answer to question 1.3/1.4. Don't forget that if you are using an a.out system you must first have installed the dld package (which is not supplied with most a.out systems) to have dlfcn.h at all. See Question 1.11.

37.6 GCC complains about an ignored option -fpic

Earlier versions of GCC accepted either -fpic or -fPIC. It appears that more recent versions (V2.7.2?) require -fPIC. If you are using an ELF version of Linux, this can safely be ignored as -fPIC is the default.

You can correct this by editing src/Makefile.global and changing CFLAGS_SL

37.7 I get warnings of the form warning- cast from pointer to integer of different size

These were seen in earlier versions of Postgres95 and could safely be ignored. PostgreSQL V6.0 should compile with no warnings except those related to system header files (which can also be safely ignored).

37.8 $SuSE - Linux4.2$ Where is curses and termcap?

SuSE-Linux has ncurses but not curses. Set the value of CURSES_LIB in src/Makefile.custom to -lncurses (or do this through the customize script).

SuSE-Linux has the termcap library is in /usr/lib/termcap instead of in /usr/lib. If you have a problem, you need to add the following line to src/Makefile.custom:

LDADD_BE+= -L/usr/lib/termcap

You may need to edit src/bin/psql/Makefile and comment out the change: ifeq $((PORTNAME), linux) LD\_ADD+ = -ltermcapto : ifeq$((PORTNAME), linux) LD_ADD+=

37.9 Why do I get problems with ld.so?

If you get problems with ld.so, another library required under ELF for dynamic loading, then you have messed up your installation or (more likely) upgrade of Linux.

See the answers to Question 1.3/1.4. You may need to install ld.so.x.y.z in /lib and run ldconfig.

The most recent stable release of the ld package is 1.7.14 At the time of writing, 1.8.x versions of ld are experimental.

37.10 Why do I get 'yy_flush_buffer undefined' errors?

This isn't really Linux specific, but is common on older Linux installations. You must have a recent version of flex (2.5.2 or later) to compile PostgreSQL. Note that flex 2.5.3 has a bug: see Question 3.4.

37.11 How do I compile PostgreSQL on an a.out system?

First, you must install the dld library. This may be obtained from Sunsite as: Linux/libs/dld.3.2.7.tar.gz

Second, add the following line to src/Makefile.custom: LINUX_ELF= (or use the customize script)

37.12 Why does make fail with - yacc -d /disk2/PostgreSQL/src/backend/parser/gram.y make- /usr/bin/make- cannot execute binary file

This was a problem in earlier versions of Postgres95. The default for PostgreSQL is to use bison -y rather than yacc.

yacc is generally implemented as a script which invokes bison -y For some reason (certain versions of make? certain versions of bash?) make is unable to execute this script file.

To correct this, simply edit src/mk/port/postgres.mk.linux and, at the end of the file, change: # YACC = bison -y to YACC = bison -y

37.13 What are the references in X11_LIB to libsocket and libnsl in src/Makefile.global?

This was a problem in 1.08 (they are Sun Solaris specific). It is fixed in 1.09 and 6.0

37.14 *DEBIAN* Where is libtermcap?

Debian Linux comes without the termcap library and uses ncurses (which uses terminfo instead). There is no need to change the CURSES_LIB variable in src/bin/psql/Makefile since Debian provides a link from libncurses to libcurses (unlike SuSE-Linux — see Question 1.8).

You may need to edit src/bin/psql/Makefile and comment out the change: ifeq $((PORTNAME), linux) LD\_ADD+ = -ltermcapto : ifeq($(PORTNAME), linux) LD_ADD+=

38 FAQ Linux-PostgreSQL - Compiling accessory programs

38.1 The linker fails to find libX11 when compiling pgtclsh

Add the following to src/Makefile.custom X11_LIBDIR = /usr/X11R6/lib

39 FAQ Linux-PostgreSQL - Runtime Problems

39.1 I get an error reporting _fUnKy_POSTPORT_sTuFf_ undefined when running scripts like createuser

This is a bug in V1.06-V1.07 of Postgres and is fixed in V1.08 and above.

39.2 I run postmaster and after that system says 'Bad system call(Core dumped)'

This indicates that you have not compiled shared memory support into your kernel. You need to recompile the Linux kernel to add this feature.

39.3 When I try to start the Postmaster, why do I get an error of the form Failed Assertion("!(file != 0) - (null)", File

Failed Assertion("!(file != 0):(null)", File: "/usr/local/PostgreSQL/src/backend/storage/file/fd.c", Line: 257) !(file != 0) (0) initdb: could not create template database initdb: cleaning up.

Your permissions on the file /dev/null are wrong.

ls -l /dev/null should give you something like:

crw-rw-rw- 1 root wheel 2, 2 Oct 8 18:41 /dev/null

Correct the permissions using:

chmod a+rw /dev/null

39.4 Why doesn't createuser work?

There is a problem with Version 2.5.3 of GNU flex and createuser. Your options are to downgrade flex to V2.5.2, apply a patch to V2.5.3 (supplied in doc/README.flex) or wait for V2.5.4 which will fix the bug.

39.5 Why do I get an error like- IpcMemoryCreate- memKey=155356396 , size=760632 ,

permission=384IpcMemoryCreate: shmget(..., create, ...) failed: Invalid argument

You haven't build IPC support into your Linux kernel. You will have to rebuild the kernel and switch on this option.

39.6 Why does psql fail with- psql- can't load library 'libpq.so.1'

Psql has been compiled to link dynamically with the libpq library.

To solve this, you should log in as root and edit the file /etc/ld.so.conf Add a single line at the end which gives the name of the PostgreSQL library directory (the lib subdirectory of your PostgreSQL installation) and run /sbin/ldconfig

40 FAQ IRIX-PostgreSQL - Compiling PostgreSQL

== Frequently Asked Questions (FAQ) for PostgresSQL V6.1 IRIX Specific TO BE READ IN CONJUNCTION WITH THE NORMAL FAQ == last updated: Fri Jun 13 09:54:00 BST 1997

current maintainer: Andrew C.R. Martin (martin@biochem.ucl.ac.uk) original author: Andrew C.R. Martin (martin@biochem.ucl.ac.uk)

Changes in this version (* = modified, + = new, - = removed): +1.5) Can I install PostgreSQL under Irix 6.4?

This file is divided approximately as follows:

- 1.*) Installing Postgres95
- 2.*) Uninstalling Postgres95
- 3.*) Extending Postgres95

41 FAQ-IRIX Installing Postgres95

41.1 What extra items do I need to install Postgres95 under Irix?

You *must* have the following installed:

- a) Gnu make (installed as gmake)

You are recommended to install the following:

- a) GNU install (installed as ginstall)

You may choose to install the following:

- a) GNU readline library (if you wish psql to have readline support).
- b) tcl/tk (if you wish to compile pgtclsh)

41.2 What changes do I need to make to src/Makefile.global or src/Makefile.custom?

The easiest way to do this is to use the customize script in the src directory.

You *must* set the following variables: PORTNAME= irix5

You will also need to change the following to match your own installation: POSTGRESDIR

If you switch on the USE_TCL option, you will need to set these: TCL_INCDIR= TCL_LIBDIR= TCL_LIB = TK_INCDIR= TK_LIBDIR= TK_LIB =

You may also make any other changes you need as documented in the INSTALL file and in Makefile.global

41.3 What are the references in X11_LIB to libsocket and libnsl in src/Makefile.global?

This was a problem in 1.08 (they are Sun Solaris specific). It is fixed in 1.09 and above.

41.4 Are there any other changes I should make?

If you have installed the GNU install program (ginstall), you should add the following line to src/Makefile.custom: CUSTOM_INSTALL=ginstall

For an explanation as to why this is a good idea, see Question 2.1

41.5 Can I install PostgreSQL under Irix 6.4?

Irix 6.4 has a bug in ld which mishandles the addresses of static procedures when object files are assembled into larger object files using 'ld -r'. This bug has been reported to Silicon Graphics. The following patch should be applied as a workaround. (Supplied by Bob Bruccoleri bruc@bms.com)

*** ./backend/Makefile.orig Thu May 22 00:00:15 1997 — ./backend/Makefile Thu Jun 5 16:47:27 1997 *************** *** 54,60 **** all: postgres $ (POSTGRES_IMP) global1.bki.source local1_template1.bki.source

postgres: $ (OBJS) ../utils/version.o ! $ (CC) -o postgres $ (OBJS) ../utils/version.o $ (LDFLAGS)

$ (OBJS): $ (DIRS:%=%.dir)

— 54,64 —- all: postgres $ (POSTGRES_IMP) global1.bki.source local1_template1.bki.source

postgres: $ (OBJS) ../utils/version.o

```
! #      \$ (CC) -o postgres  \$ (OBJS) ../utils/version.o  \$ (LDFLAGS)
!       -rm -f *.o
!       find . -name "*.o" -exec cp (flower-brackets) . \;
```

! rm -f SUBSYS.o ! $ (CC) -o postgres *.o ../utils/version.o $ (LDFLAGS)

$ (OBJS): $ (DIRS:%=%.dir)

42 FAQ-IRIX Deinstalling Postgres95

42.1 Why can't I move the executable files?

By default, the IRIX port uses the BSD compatible version of install from /usr/bin/X11. If you read the man page for this version of install, you will see that it is not meant for end-user use; it has the interesting side-effect of chowning files it installs to root.

You should still be able to delete the files as you (the postgres user) will own the directory in which they are stored.

The normal IRIX install program cannot be used easily as it takes its arguments in the reverse order. It is therefore recommended to use the GNU version of install (ginstall). See Question 1.4

43 FAQ-IRIX Extending Postgres95

43.1 How do I compile a C program to create a function for extending Postgres95

Here is a sample command line:

cc -I/usr/local/postgres95/include/ -I/usr/local/postgres95/src/backend -shared -o funcs.so funcs.c

44 FAQ-IRIX PostgreSQL AUTHOR contacts

Dr. Andrew C.R. Martin University College London

EMAIL: (Work) martin@biochem.ucl.ac.uk

(Home) andrew@stagleys.demon.co.uk

URL: `<http://www.biochem.ucl.ac.uk/ martin>`

Tel: (Work) +44(0)171 419 3890

Tel: (Home) +44(0)1372 275775

45 Copyright Notice